

Release Bulletin Adaptive Server® Enterprise Version 12.5.1 for Mac OS X 10.3

Document ID: DC79176-01-1251-01

Last revised: November 25, 2003

Topic	Page
1. Accessing current release bulletin information	2
2. Product summary	3
2.1 Installation kit	3
2.2 Server components	3
2.3 PC-client components	4
2.4 Operating system updates	5
3. Special installation instructions	5
4. Special upgrade instructions	5
5. Product and platform interoperability	5
6. New functionality in this version	6
6.1 TCP_NODELAY is turned on by default	7
6.2 PowerTransfer	7
7. Known problems	7
7.2 alter table and triggers with if update() clauses	7
7.3 Truncation error using select with power() function and numeric argument	8
7.4 Highlighted known ASE Replicator problems	8
7.5 SQL derived tables cannot have unnamed columns	13
7.6 Reconfiguring from utf8 with external sort order	13
7.7 Queries containing a select statement with a group by statement	14
7.8 SybMigrate error messages	14

Copyright 1989-2004 by Sybase, Inc. All rights reserved. Sybase, the Sybase logo, AccellaTrade, ADA Workbench, Adaptable Windowing Environment, Adaptive Component Architecture, Adaptive Server, Adaptive Server Anywhere, Adaptive Server Enterprise, Adaptive Server Enterprise Monitor, Adaptive Server Enterprise Replication, Adaptive Server Everywhere, Adaptive Server IQ, Adaptive Warehouse, Anywhere Studio, Application Manager, AppModeler, APT Workbench, APT-Build, APT-Edit, APT-Execute, APT-FORMS, APT-Translator, APT-Library, AvantGo, AvantGo Application Alerts, AvantGo Mobile Delivery, AvantGo Mobile Document Viewer, AvantGo Mobile Inspection, AvantGo Mobile Marketing Channel, AvantGo Mobile Sales, AvantGo Pylon, AvantGo Pylon Application Server, AvantGo Pylon Conduit, AvantGo Pylon PIM Server, AvantGo Pylon Pro, Backup Server, BizTracker, ClearConnect, Client-Library, Client Services, Convoys/DM, Copernicus, Data Pipeline, Data Workbench, DataArchitect, Database Analyzer, DataExpress, DataServer, DataWindow, DB-Library, dbQueue, Developers Workbench, Direct Connect Anywhere, DirectConnect, Distribution Director, e-ADK, E-Anywhere, e-Biz Integrator, E-Whatever, EC Gateway, ECMAPI, ECRTP, eFulfillment Accelerator, Embedded SQL, EMS, Enterprise Application Studio, Enterprise Client/Server, Enterprise Connect, Enterprise Data Studio, Enterprise Manager, Enterprise SQL Server Manager, Enterprise Work Architecture, Enterprise Work Designer, Enterprise Work Modeler, eProcurement Accelerator, EWA, Financial Fusion, Financial Fusion Server, Gateway Manager, GlobalFIX, ImpactNow, Industry Warehouse Studio, InfoMaker, Information Anywhere, Information Everywhere, InformationConnect, InternetBuilder, iScript, Jaguar CTS, jConnect for JDBC, Mail Anywhere Studio, MainframeConnect, Maintenance Express, Manage Anywhere Studio, M-Business Channel, M-Business Network, M-Business Server, MDI Access Server, MDI Database Gateway, media.splash, MetaWorks, My AvantGo, My AvantGo Media Channel, My AvantGo Mobile Marketing, MySupport, Net-Gateway, Net-Library, New Era of Networks, ObjectConnect, ObjectCycle, OmniConnect, OmniSQL Access Module, OmniSQL Toolkit, Open Biz, Open Client, Open Client/Server, Open Client/Server Interfaces, Open Gateway, Open Server, Open ServerConnect, Open Solutions, Optima++, Orchestration Studio, PB-Gen, PC APT Execute, PC Net Library, PocketBuilder, Pocket PowerBuilder, Power++, powerstop, PowerAMC, PowerBuilder, PowerBuilder Foundation Class Library, PowerDesigner, PowerDimensions, PowerDynamo, PowerJ, PowerScript, PowerSite, PowerSocket, Powersoft, PowerStage, PowerStudio, PowerTips, Powersoft Portfolio, Powersoft Professional, PowerWare Desktop, PowerWare Enterprise, ProcessAnalyst, Rapport, Report Workbench, Report-Execute, Replication Agent, Replication Driver, Replication Server, Replication Server Manager, Replication Toolkit, Resource Manager, RW-DisplayLib, S-Designer, SDF, Secure SQL Server, Secure SQL Toolset, Security Guardian, SKILLS, smart.partners, smart.script, SQL Advantage, SQL Anywhere, SQL Anywhere Studio, SQL Code Checker, SQL Debug, SQL Edit, SQL Edit/TPU, SQL Everywhere, SQL Modeler, SQL Remote, SQL Server, SQL Server Manager, SQL SMART, SQL Toolset, SQL Server/CFT, SQL Server/DBM, SQL Server SNMP SubAgent, SQL Station, SQLJ, STEP, SupportNow, SWIFT, Message Format Libraries, Sybase Central, Sybase Client/Server Interfaces, Sybase Financial Server, Sybase Gateways, Sybase MPP, Sybase SQL Desktop, Sybase SQL Lifecycle, Sybase SQL Workgroup, Sybase User Workbench, SybaseWare, Syber Financial, SyberAssist, SyBooks, System 10, System 11, System XI (logo), SystemTools, Tabular Data Stream, TotalFix, TradeForce, Transact-SQL, Translation Toolkit, UltraLite.NET, UNIBOM, Unidb, Uninall, Unisep, Unistring, URK Runtime Kit for UniCode, Viewer, Visual Components, VisualSpeller, VisualWriter, VQL, WarehouseArchitect, Warehouse Control Center, Warehouse Studio, Warehouse WORKS, Watcom, Watcom SQL, Watcom SQL Server, Web Deployment Kit, Web.PB, Web.SQL, WebSigns, WebViewer, WorkGroup SQL Server, XA-Library, XA-Server and XP Server are trademarks of Sybase, Inc. 07/03

Topic	Page
7.9 Errors in queries with derived tables	14
7.10 Views must have defined column names	14
7.11 Migrating between single-byte and multibyte charactersets	15
7.12 cis connect timeout and enable SNMP do not work	15
7.13 XP Server returns incorrect results	15
7.14 Network license daemon support	15
7.15 bcp and abstract data types	16
7.16 Referencing Java data in a nonrecovered database	16
7.17 Configuring connectivity for clients accessing wide data	16
7.18 Handling multibyte character sets during migration	16
7.19 Adaptive Server plug-in default character set and language	17
7.20 Dynamic and host variables not allowed in DDL statements	17
8. Product compatibilities	17
8.1 Replication Server compatibility	18
9. Documentation updates and clarifications	19
9.1 Utility Guide	19
9.2 Transact-SQL User's Guide	22
9.3 Reference Manual: Building Blocks	22
9.4 Reference Manual: Commands	22
9.5 System Administration Guide	25
9.6 Performance and Tuning Guide	26
10. Technical support	27
11. Other sources of information	27
11.1 Sybase certifications on the Web	28
11.2 Sybase EBFs and software maintenance	28

1. Accessing current release bulletin information

A more recent version of this release bulletin may be available on the Web. To check for critical product or document information added after the release of the product CD, use the Sybase Technical Library Product Manuals Web site.

v Accessing release bulletins at the Technical Library Product Manuals Web site

- 1 Go to Product Manuals at <http://www.sybase.com/support/manuals/>.
- 2 Follow the links to the appropriate Sybase product.
- 3 Select the Release Bulletins link.

- 4 Select the Sybase product version from the Release Bulletins list.
- 5 From the list of individual documents, select the link to the release bulletin for your platform. You can either download the PDF version or browse the document online.

2. Product summary

Enclosed is Sybase® Adaptive Server® Enterprise version 12.5.1. Server and client components are distributed on separate CDs.

For details on system requirements, including disk space and RAM, see the installation guide for your platform.

For more information specific to jConnect™ for JDBC™, jisql, and Ribo, see the jConnect documentation.

This release bulletin provides the latest information about Adaptive Server and its known problems, as well as documentation updates and clarifications.

2.1 Installation kit

The installation kit includes:

- The server CD – for contents list, see “Server components” on page 3.
- The PC-client CD – contains software client components to be installed on Windows NT, Windows 2000, XP Pro, and Windows 2003 computers.
- *Release Bulletin for Adaptive Server Enterprise Version 12.5.1 for Mac OS X 10.3*(this document).
- *Quick Installation Guide for Sybase Adaptive Server Enterprise Version 12.5.1 for Mac OS X 10.3*.

2.2 Server components

The Server CD includes the following components:

- Adaptive Server 12.5.1
- Language Modules 12.5.1
- jConnect for JDBC 5.5
- Java utilities:

- jisql 2.0
- Ribo 2.0
- Sybase Central™ 4.x
- Server Admin plug-in
- System Preferences panel
- Server Discovery
- ODBC driver
- SNMP sub-agent

2.3 PC-client components

The PC-client CD includes the following components, which can be installed on Windows:

- Software Development Kit (SDK) 12.5.1 ASE Edition, which includes the following:
 - Open Client 12.5.1 ASE Edition
 - Embedded SQL/C 12.5.1 ASE Edition
 - Embedded SQL/Cobol 12.5.1 ASE Edition
- Open Client Runtime 12.5.1
- ODBC 4.2
- OLE DB 2.7
- jConnect 5.5 and related documentation
- Java utilities:
 - jisql 2.0
 - Ribo 2.0
- InfoMaker® 9.0
- Language Modules 12.5.1
- XA Interface Library for ASE Distributed Data Manager 12.5.1
- SQL Remote 7.0.2
- Monitor Client Library 12.5.1

- PowerDesigner 9.5.2
- PowerTransfer 9.5.2

2.4 Operating system updates

Adaptive Server version 12.5.1 requires that you run Mac OS X 10.3 or later.

Viewing installed patches

On Mac OS X 10.3, to view the current version of the operating system, go to the Apple menu button in the upper left corner, and select About this Mac. A window displays, and shows the current version of the operating system.

3. Special installation instructions

There are no known installation issues for Mac OS X 10.3 on Adaptive Server version 12.5.1.

4. Special upgrade instructions

This section provides information for upgrading to Adaptive Server version 12.5.1 that is not included or should be corrected in the installation guide.

4.1 Upgrading from Jaguar to Panther

If you intend to upgrade your Mac OS X system from version 10.2.x to 10.3, Adaptive Server version 12.5.0.3 and earlier will not work on Mac OS X 10.3. After upgrading to Mac OS X 10.3, Sybase recommends that you upgrade your Adaptive Server to version 12.5.1

5. Product and platform interoperability

This section shows the interoperability of Adaptive Server, Replication Server, and Open Client/Server (OCS) products across versions. For specific platform or operating system level information, see the respective product certification reports at <http://www.sybase.com/>.

Table 1: Interoperability between Replication Server, Open Client/Server, and Adaptive Server

Adaptive Server 12.5.1	Adaptive Server 12.0	Open Client/Server 12.5.1	Open Client/Server 12.0	Replication Server 12.6	Replication Server 12.5	Replication Server 12.1
Apple Mac OS X 10.3	N/A	N/A	N/A	X	N/A	N/A
HP Tru64 UNIX 5.x	X	X	X	X	X	X
HP HP-UX 11.x	X	X	X	X	X	X
IBM RS 6000 AIX 4.3.x	X	X	X	X	X	X
IBM RS 6000 5.1	X	X	X	X	X	X
Red Hat Enterprise 2.1	N/A	X	N/A	X	X	N/A
Silicon Graphics IRIX	N/A	X	N/A	X	X	N/A
Sun Solaris	X	X	X	X	X	X
Microsoft Windows	X	X	X	X	X	X

Legend

- X = compatibility
- N/A = incompatibility, or the product is not available for that version/platform combination

Note 1 12.5.1 is a shelf-replacement for the 12.5 versions of Adaptive Server and Open Client/Server. Therefore, the 12.5 version of either product is compatible with 12.5.1, although Sybase strongly recommends that customers upgrade to 12.5.1 as soon as possible.

Note 2 Even though two or more products may be interoperable, features introduced in a newer version of a product are not likely to be supported with older versions of the same or other products.

Note 3 There is no 12.5.1 version of SDK or Open Server on Mac OS X 10.3. The 12.5 SDK and Open Server will be certified on Mac OS X 10.3.

Note 4 While Replication Server 12.5 is not available on the Mac OS X 10.3 platform Adaptive Server on Mac OS X 10.3 is compatible with Replication Server 12.5 that is running on other platforms.

For the latest interoperability report on these products, see the MySybase Web site at <http://my.sybase.com/detail?id=1026087>

6. New functionality in this version

This section describes changed functionality in Adaptive Server version 12.5.1 that is not included in the *What's New in Adaptive Server 12.5.1*.

New features specific to Adaptive Server on Mac OS X 10.3 are discussed in the *Quick Installation Guide*.

6.1 TCP_NODELAY is turned on by default

TCP_NODELAY is turned on by default, with changes made by the user to the configuration file.

6.2 PowerTransfer

PowerTransfer is a PC Client component and is an extension to PowerDesigner that allows the transfer of data into Adaptive Server. The source database can be any ODBC data source supported by PowerDesigner. PowerTransfer uses the Sybase bulk copy mechanism (bcp) for high-performance inserts into Adaptive Server.

7. Known problems

This section describes known problems and workarounds for Adaptive Server.

7.1 Use of worker threads with multiple engines

[CR #339128] If you are using multiple engines on a dual G5 system, there is a problem with worker threads that causes fatal statement errors.

Workaround: Do not configure worker threads if using dual G5 systems and Adaptive Server is configured to use more than one engine. If worker threads are necessary, configure Adaptive Server for single engine usage.

7.2 *alter table* and triggers with *if update()* clauses

[CR #199655] If a trigger contains an *if update()* clause, data modifications that should fire the trigger that are performed after executing *alter table add*, *alter table drop*, *alter table lock*, or *alter table modify* may cause errors in column references. Triggers on the altered table that use an *if update()* clause in the body of the trigger to reference a column may not fire, or may fire incorrectly.

Workaround: After the *alter table* operation has completed, drop and re-create all triggers on the altered table. Doing so causes the *if update()* clause in the triggers to correctly reference the new columns by their new column offsets, so the trigger code executes correctly.

7.3 Truncation error using *select* with *power()* function and numeric argument

[CR #236723] If you use the *select* command with the function *power()*, and use a numeric argument, the result is a truncation error.

Workaround: Convert the numeric argument to a float.

7.4 Highlighted known ASE Replicator problems

This section describes known problems and workarounds for ASE Replicator.

7.4.1 Subscription status incorrect after abnormal shutdown

[CR #270782] The subscription status returned by *sp_helprep* and *sp_helpsub* is incorrect under certain conditions.

This problem occurs after the ASE Replicator process is shut down abnormally, such as by a power failure on the host machine or a process kill command, while a subscription is resumed (up), and the remote server associated with that subscription is still down when ASE Replicator is restarted. When ASE Replicator is restarted, the subscription status is actually 0 (down), but *sp_helprep* and *sp_helpsub* return the subscription status as 1 (up).

Workaround: This is an error only in the status information returned. After the remote server has successfully restarted, resume the subscription normally by using *sp_resumesub*.

7.4.2 No errors with invalid database connection name

[CR #s 270729, 271653, and 271676] The following ASE Replicator procedures do not return an error message when you invoke them with an invalid or null database connection name:

- *sp_helprep*
- *sp_resumeprimaryconn*
- *sp_resumereplicateconn*
- *sp_suspendprimaryconn*
- *sp_suspendreplicateconn*

Workaround: There is no workaround for this issue.

7.4.3 Incorrect error message with invalid database connection name

[CR #269484] The following ASE Replicator procedures return an incorrect error message when you invoke them with an invalid or null database connection name, or an invalid or null connection configuration parameter name:

- `sp_configprimaryconn`
- `sp_configreplicateconn`

Workaround: There is no workaround for this issue.

7.4.4 Incorrect error message with invalid subscription name

[CR #271677] When you invoke `sp_helprep` with the status keyword and an invalid subscription name (`sub=sub_name`), it returns an incorrect error message.

Workaround: There is no workaround for this issue.

7.4.5 Statistics not returned by some commands

[CR #270080] The following ASE Replicator procedures do not return a statistics result set:

- `sp_helprep stats`
- `sp_helprep stats, pri_conns`
- `sp_helprep stats, subs`
- `sp_helprep stats, sub=subname`
- `sp_helpreplicateconn stats`
- `sp_helpsub stats`

Workaround: To get statistics for a primary connection, use the following command:

```
sp_helpprimaryconn stats
```

To get statistics for a replicate connection, use the following command:

```
sp_helprep stats, rep_conns
```

7.4.6 UNIX permission problems with RUN script creation

[CR #279616] When *aserep.sh* creates the *RUN* script, it sets permissions on the *RUN* script file to allow execution. If *aserep.sh* encounters a problem setting permissions on the *RUN* script file, the following error message is returned:

```
Component message: Problem setting permissions;  
exitValue = 1
```

Workaround: To solve this problem:

- 1 Log in to the operating system with a user ID that has authority to set execute permissions in the ASE Replicator instance subdirectory.
- 2 Execute the *aserep.sh* script to initialize the ASE Replicator instance.

See Chapter 2, “Setting Up and Starting ASE Replicator,” in the *ASE Replicator User’s Guide* for more information on executing the *aserep.sh* script to initialize the ASE Replicator instance.

7.4.7 Publishing multiple tables with the same name not supported

[CR #279148] Adaptive Server allows non-unique table names in a database, so a specific table must be identified by a qualified object name. ASE Replicator does not fully support qualified object names when identifying primary tables to be published.

If you use owner-qualification in `sp_addprimaryart`, the correct table (identified by its qualified object name) is published, but you cannot publish additional tables that have the same name in that database.

Workaround: There is no workaround for this issue.

7.4.8 Some nested quotes not allowed

[CR #269224] When parameters passed to ASE Replicator command procedures contain nested quotes, outer quote characters must be double quotes and inner (nested) quote characters must be single quotes. For example:

```
sp_addreplicateart $subname, $pubobj, $subobj,  
@where_clause = "where col1 = '3' "
```

Workaround: Always use double quote characters outside single quote characters when you use strings with nested quotes in ASE Replicator command procedure parameters.

7.4.9 Java object replication not supported

[CR #s 270775 and 286817] ASE Replicator does not support replicating Java objects.

Workaround: There is no workaround for this issue.

7.4.10 Some commands fail on case-insensitive servers

[CR #270776] ASE Replicator command procedures that deal with articles may fail on case-insensitive servers.

Workaround: Use the exact case of the database object name returned by catalog stored procedures on the case-insensitive server.

7.4.11 Some messages returned incorrectly on UTF-8 servers

[CR #271646] When the Adaptive Server default character set is UTF-8, the ASE Replicator procedures `sp_addreplicatart` and `sp_addreplicateconn` return some of their error messages incorrectly. This problem does not affect the function of these procedures.

Note Although the error messages are returned incorrectly, the errors are logged correctly in the ASE Replicator log files.

Workaround: There is no workaround for this issue.

7.4.12 Roman-8 character set issues

[CR #291478] The ASE Replicator process uses `jConnect` for JDBC for all of its communication with Adaptive Server. Therefore, the Adaptive Server default character set must map to a character set supported by `jConnect`. If the Adaptive Server default character set does not map to a character set supported by `jConnect` (such as Roman-8, for example), ASE Replicator cannot connect to the Adaptive Server.

When this problem occurs, the following error message appears in the `system.log` file:

```
Message: JZ0IB: The server's default charset of roman8
does not map to an encoding that is available in the
client Java environment.
```

Workaround: Use an Adaptive Server default character set that maps to a character set supported by `jConnect` (such as UTF-8, for example).

7.4.13 Java out-of-memory problems

[CR #273087] In some circumstances, ASE Replicator may shut down with a Java `OutOfMemory` error. This problem can occur when:

- The value of the `queue_size` parameter is set too high for either a database connection or the ASE Replicator general configuration parameter. Sybase recommends that you use the default values.
- A problem occurs on one or more connections where data gets backed up in the queue. Depending on the size of the data, the `queue_size` value for the affected connection, and the amount of memory allocated to the Java virtual machine (VM), a Java `OutOfMemory` exception can occur.

Workaround: If you encounter a Java `OutOfMemory` error, you can do either or both of the following:

- Reduce the `queue_size` for each database connection and for the ASE Replicator general configuration.
- Increase the amount of memory allocated to the Java VM by editing the `-Xmx` parameter in the `aserep` start-up script. The default is 64MB.

You can estimate the amount of memory to allocate to the Java VM for a particular `queue_size` value by evaluating the average size of a row of data in any table at each database connection, and factoring that value with the amount of memory available on the machine.

- For each primary database connection:
 - Determine the average size (in bytes) of a row of table data replicated.
 - Multiply that value by the `queue_size` value for the connection, and add an overhead factor of 25 percent of the data size.
 - Multiply that number by the number of subscriptions that subscribe to data in a publication on the primary connection.
- Find the average (or maximum) for all primary connections and multiply that number by the ASE Replicator `queue_size` setting. Add an overhead factor of 25 percent.
- The sum of these numbers, plus approximately 32MB, should give you an estimate of the amount of memory required by the Java VM.

Note This amount of memory is required only in a worst-case scenario, in which the queues are filled with data.

7.5 SQL derived tables cannot have unnamed columns

[CR #333873] SQL derived tables cannot have unnamed columns.

For example:

```
select * from (select sum(advance) from titles) d1
```

causes an error message.

Workaround: Rewrite the query using derived column list or the “AS” clause.

For example:

```
select * from (select sum(advance) from titles) d1(a1)
```

or:

```
select * from (select sum(advance) as a1 from titles) d1
```

7.6 Reconfiguring from utf8 with external sort order

[CR #333614] Before Adaptive Server version 12.5.1, the utf8 character set could be used only with the binary sort order. However, Adaptive Server release 12.5.1 supports non-binary sort orders with the utf8 character set.

If you use utf8 and then configure an external sort order (see Table 2), you cannot reconfigure Adaptive Server with another character set without first reverting to the binary sort order with utf8. After reverting to binary sort order (which has an ID of 50), you can then safely change to any other character set.

If you change the character set without first going to the binary sort order with utf8, recovery of the master database hangs and Adaptive Server displays error message 1569:

```
Unable to get a conversion handler for the built-in sort function.
```

If this occurs, you must download and start Adaptive Server version 12.5.1 ESD #1 or later to successfully bring up a server (see the Sybase Web site at <http://www.sybase.com> for information about downloading Adaptive Server ESDs). Check with Sybase support for the availability of Adaptive Server version 12.5.1 ESD #1.

Alternatively, you can re-create the master device and load from database dumps.

Table 2 lists the external sort orders for utf8 and their IDs.

Table 2: External sort orders

Sort order	Sort order ID
big5bin	194
cp932bin	192
cyrdict	140
eucjisbn	192
euckscbn	161
gb2312bn	137
gbpinyinbn	163
rusdict	165
sjisbin	179

`$SYBASE/collate/unicode` also lists the sort orders and their IDs

7.7 Queries containing a select statement with a group by statement

[CR #333077] Queries containing a select with a group by with a case expression using a local variable, may return wrong results.

Workaround: Remove the use of the local variable in group by.

7.8 SybMigrate error messages

[CR #332051] When starting SybMigrate, warning messages display in the log.

Workaround: These messages are harmless and can be ignored.

7.9 Errors in queries with derived tables

[CR #329454] Incorrect results may be returned for a query having more than one derived table with group by aggregates in the derived table expression.

Workaround: Do not run these types of queries.

7.10 Views must have defined column names

[CR #314976] ANSI requirements specify that a view must have defined column names.

Workaround: You can suppress this requirement by using trace flag 1702.

7.11 Migrating between single-byte and multibyte charactersets

[CR #295842] SybMigrate does not support migration of data from single-byte (iso_1, roman8) to multibyte (utf8), even with -Jutf8.

Workaround: There is no workaround for this issue.

7.12 *cis connect timeout* and *enable SNMP* do not work

[CR #323177] Configuration parameters *cis connect timeout* and *enable SNMP* are not currently implemented.

Workaround: Do not use these configuration parameters.

7.13 XP Server returns incorrect results

[CR #280073] On all systems, using an extended stored procedure (ESP) with parameters may result in wrong data being sent to XP Server. Because of this, XP Server may return incorrect results.

This happens when there is no parameter declaration list in the ESP, but one or more parameters are supplied at execution time. Before CR #280073, all parameters supplied at execution time were sent to the XP Server, causing incorrect results. CR #280073 fixes this problem by silently ignoring all parameters supplied at execution time when the parameter declaration list is empty.

This fix does not cause any behavior change when the parameter declaration list in the ESP matches the parameters supplied at execution time.

If you are using ESPs without a parameter declaration list, and you supply parameters at execution time, drop and re-create the ESP in Adaptive Server. Make sure the parameter declaration list in the ESP matches the parameter list supplied at execution time

You can use trace flag 523 to revert to pre-CR280073 behavior. This is a server-wide trace flag to be set when you restart Adaptive Server.

7.14 Network license daemon support

Problem: Adaptive Server version 12.5.1 has upgraded the SySAM daemons and libraries (FlexLM) to version 8. The SySAM daemons (Imgrd and SYBASE) are backward compatible, however the libraries are not. As a result, Adaptive Server 12.5.1 cannot check out licenses from earlier versions of the SySAM license daemons.

Workaround: Sites using the Network License Model should make sure that the central SySAM license daemons are upgraded first.

7.15 *bcp* and abstract data types

[CR #268302] Adaptive Server allows you to use *bcp* in to copy data containing Java objects, but it performs no load-time checking to ensure that the objects are consistent with the classes that are installed in the server. Any checking is delayed until the objects are actually accessed by an application. To be sure that all objects are compatible with their corresponding classes, Sybase recommends that you run a *select* query to access these objects on the server immediately after they have been copied. If there are any incompatibilities, they are reported via a Java serialization error (10707).

Workaround: There is no workaround for this issue.

7.16 Referencing Java data in a nonrecovered database

[CR #274071] Referencing Java data in a database that is not recovered or operational causes a “*process (0xf000f) infected with 11*” condition, which terminates the database connection.

Workaround: Verify that the database is recovered and available for use by issuing:

```
use database database_name
```

7.17 Configuring connectivity for clients accessing wide data

[CR #271822] If you are using wide data (that is, data other than text and image that is greater than 255 bytes), and you are using the bulk library APIs, make sure you configure the connectivity capability for wide data. Failure to do so may result in memory corruption.

7.18 Handling multibyte character sets during migration

[CR #271839] If you configure multibyte character sets after migrating data from system catalogs *sysattributes* and *sysxtypes*, the text columns in these catalogs are inconsistent with the multibyte character sets.

Workaround: Manually run *dbcc fix_text* on *sysattributes* and *sysxtypes*, to make the text columns consistent with the multibyte character sets.

7.19 Adaptive Server plug-in default character set and language

[CR #270906] The Adaptive Server plug-in default character set is iso_1, and the default language is us_english. If you are running an Adaptive Server that is using a different language or character set, you may encounter compatibility problems.

Workaround: Set the character set and language values in the Adaptive Server plug-in Connect dialog, so that they correspond to the configuration of the Adaptive Server.

7.20 Dynamic and host variables not allowed in DDL statements

[CR #234217] Dynamic and host variables in create view statements were allowed in earlier versions, but worked incorrectly when:

- create view statements were replicated, or
- Views containing dynamic and host variables were upgraded from text in the syscomments table.

Dynamic and host variables are no longer allowed in create view statements.

Workaround: To allow dynamic and host variables with this incorrect behavior in create view statements, start Adaptive Server with trace flag 108:

```
dataserver -dd_master -T108
```

8. Product compatibilities

This section lists the Sybase components that are compatible with Adaptive Server Enterprise version 12.5 and later. For information about operating system requirements, see individual component documentation.

The following components have been tested for compatibility with Adaptive Server 12.5 and later:

- Sybase Character Sets 3.0
- Sybase Central viewer 4.x
- ECDA option for ODBC 12.5
- ECDA option for Informix 12.5
- MainframeConnect DirectConnect for OS/390 12.5
- ECDA option for Oracle 12.5

- ECDA option for Microsoft SQL Server 12.5
- ECDA option for DB2 Universal Database 12.5
- ECDA option for AS/400 12.5
- OpenSwitch 12.0
- jConnect for JDBC 4.5
- jConnect for JDBC 5.5
- Open Server™ 12.5
- SDK 12.5
- Open ServerConnect™ (CICS, IMS/MVS) 4.0
- Open ClientConnect™ (CICS, IMS/MVS) 4.0
- ODBC Driver for Adaptive Server 12.5
- OLE DB Provider for Adaptive Server 12.5
- InfoMaker 7.0.3
- Replication Server® 11.5.1, 12.0
- Security Guardian™ 11.1.1
- SQL Remote 6.0.2
- Physical Architect 8.0
- XA-Library for CICS/ENCINA 11.1.1
- XA-Server for Tuxedo 11.1.1

Note Backup Server 12.5 is compatible with Backup Server 11.9.2.1 and 12.0. Either server can be local or remote. Backup Server 12.5 is compatible with versions of Backup Server earlier than 11.9.2.1 *only* if the Backup Server 12.5 is the remote server.

8.1 Replication Server compatibility

Replication Server
version 12.1
compatibility

Replication Server version 12.1:

- Does not replicate extended data such as wide rows or wide columns, without truncating it.
- Does not replicate Unicode data.

- Does not use SSL when connecting to Adaptive Server, although Replication Server can still replicate data from and to an Adaptive Server that uses SSL.
- Does not use an LDAP server to retrieve server information, but instead relies on traditional interfaces files for this information. Replication Server can, however, still replicate data from and to an Adaptive Server that uses an LDAP server for this purpose.

DirectConnect and
OpenSwitch
compatibility

Direct Connect and OpenSwitch do not currently support the following Adaptive Server version 12.5 and later features:

- Extended limits, such as wide rows and wide columns
- Unicode
- SSL
- LDAP

9. Documentation updates and clarifications

This section describes changes and additions to the Adaptive Server version 12.5.1 documentation set.

9.1 Utility Guide

9.1.1 *dataserver* clarification

In the *dataserver* reference page in the *Utility Guide*, Example 2 is incorrect. It currently reads:

Example 2 rewrites a corrupt model database:

```
dataserver -d my_master_device -w model
```

The correct example is:

Example 2 rewrites a corrupt model database:

```
dataserver -d d_master -w model -s server_name
```

9.1.2 *ddlgen* update

The existing description of the `-S` parameter in *ddlgen* is:

- `-Shost_name : port_number`

specifies the host name or IP address of Adaptive Server, as well as its port number. Separate *host_name* and *port_number* with a colon, without spaces before or after it.

Note You must use the -S option because `ddlgen` does not connect to a default server.

The updated description is as follows:

- -S*server* – specifies the name of the Adaptive Server. `ddlgen` looks this name up in the interfaces file or LDAP configuration file. If you specify:
 - -S [*host_name:port_number*] – `ddlgen` uses the *host_name* and *port_number* provided, and neither interfaces nor LDAP configuration files are read or parsed.
 - -S[*server*] -l – `ddlgen` parses the interfaces file specified at the user location for the server name (see the -l parameter description for more information).
 - -S[*server*] – without specifying an interfaces file, `ddlgen`:
 - 1 Tries to read the LDAP configuration file from the standard location.
 - 2 If the LDAP file does not exist, or exists but does not contain an Adaptive Server entry, then the interfaces file is parsed at its standard location for the server name.
 - 3 If the LDAP file exists, then `ddlgen` uses it to search the server name. The interfaces file is not parsed, and the LDAP configuration file is parsed.

Separate *host_name* and *port_number* with a colon, without spaces before or after it.

Note You must use the -S option because `ddlgen` does not connect to a default server.

- -l – specifies the interfaces file name, and corresponds to `$$SYBASE/interfaces` for UNIX, and `%SYBASE%\ini\sql.ini` for Windows. Use this optional parameter with -S.

9.1.3 Omission from *SybMigrate*

The following information was omitted from the “SybMigrate” Chapter, in the *Utility Guide*.

SybMigrate requires allow resource limits to be set to 0.

Note If metadata already exists on the target Adaptive Server, you cannot migrate server data.

9.1.4 Utilities moved

The following three utilities have moved from `$$SYBASE/$$SYBASE_ASE/bin` to `$$SYBASE/$$SYBASE_OCS/bin`:

- certauth
- certpk12
- certreq

9.1.5 *bcp* clarification

In the section “bcp modes” in Chapter 3, “Using bcp to Transfer Data to and from Adaptive Server” in the *Utility Guide*, the definitions of bcp in were reversed and incorrectly stated as:

bcp in works in one of two modes:

- Fast bcp – logs each row insert that it makes, used for tables that have one or more indexes or triggers.
- Slow bcp – logs only page allocations, copying data into tables without indexes or triggers at the fastest speed possible.

The correct definitions are as follows:

bcp in works in one of two modes:

- Slow bcp – logs each row insert that it makes, used for tables that have one or more indexes or triggers.
- Fast bcp – logs only page allocations, copying data into tables without indexes or triggers at the fastest speed possible.

9.2 Transact-SQL User's Guide

In “Sorting Query results” in Chapter 3, the word “aggregates” has been dropped from this sentence:

“Adaptive Server does not allow subqueries, variables and constant expressions in the order by list.” A successful example in the following section, order by and group by, uses the aggregate function avg(price) in an order by clause:

```
order by avg(price)
```

9.3 Reference Manual: Building Blocks

9.3.1 @@version_as_integer update

In Chapter 3, “Global Variables” in *Reference Manual: Building Blocks*, the definition for @@version_as_integer states:

“Returns the version of the current release of Adaptive Server as an integer.”

This information should read:

“Returns the number of the last upgrade version of the current release of Adaptive Server as an integer. For example, @@version_as_integer returns 12500 if you are running Adaptive Server version 12.5, 12.5.0.3, or 12.5.1.”

9.3.2 set_appcontext update

In the reference page for set_appcontext in Chapter 2, “Transact-SQL Functions” in *Reference Manual: Building Blocks*, the parameter for attribute_value is incorrectly stated as being a char(2048). The correct datatype is char(30), and should appear as follows:

attribute_value – is a row that specifies an application attribute value. It is saved as the datatype char(30).

9.4 Reference Manual: Commands

9.4.1 dbcc check parameter

The following description of the dbcc check command should be included in *Reference Manual: Commands*.

Syntax	<code>dbcc upgrade_object({dbid dbname}, { compiled_object_name "view" "procedure" "default" "rule" "default" "trigger" }, ["force" "check"])</code>
Keywords and options	check – checks syntax for the specified <i>compiled_object</i> in syscomments in the specified database. Does not raise errors on occurrences of select.
Comments	<ul style="list-style-type: none"> • check is used to detect syscomments text corruption which was caused by Adaptive Server defects that occurred before the upgrade. This syscomments text corruption is serious because it causes upgrade to fail. • If any error is reported by <code>dbcc upgrade_object</code> with check option, you must drop and re-create the <i>compiled_object</i>.

9.4.2 create index clarification

The create index reference page in Chapter 1, “Commands” of *Reference Manual: Commands* incorrectly states that the `max_rows_per_page` and `fillfactor` parameters are exclusive.

“`max_rows_per_page` – `max_rows_per_page` and `fillfactor` are mutually exclusive. Unlike `fillfactor`, the `max_rows_per_page` value is maintained until it is changed with `sp_chgattribute`.”

It should read:

`max_rows_per_page` – limits the number of rows on data pages and the leaf level pages of indexes. Unlike `fillfactor`, the `max_rows_per_page` value is maintained until it is changed with `sp_chgattribute`.

9.4.3 bcp reference page clarification

In the bcp reference page in Chapter 8, “Utility Commands Reference” of the *Reference Manual: Commands*, the description of the `-b` batchsize parameter has been clarified. It used to read:

“`-b` batchsize is the number of rows per batch of data copied (default is to copy all the rows in one batch).”

The text now reads:

“`-b` batchsize is the number of rows per batch of data copied. By default, bcp in copies `n` rows in one batch, where `n` is equal to the batch size.”

9.4.4 set command option *explicit_transaction_require* not supported

The set command option *explicit_transaction_require* is not currently supported, and should not appear in the *Adaptive Server Reference Guide* version 12.5.1.

9.4.5 New set parameters *proc_output_params* and *set proc_return_status*

The following new parameters should be included in the set command reference section of *Reference Manual: Commands*.

Parameters

- *proc_output_params* – controls sending of output parameters that a stored procedure generates back to the client. *set proc_output_params off* suppresses sending the output parameters back to the client. The default for this parameter is on.
- *proc_return_status* – controls sending of a return status TDS token back to the client. *set proc_return_status off* suppresses sending the return status token to the client, and isql client does not display the (*return status* = 0) message. The default for this parameter is on.

Warning! If the client application that executes a procedure relies on the success or failure of the procedure based on the return status, then do not use the *set proc_return_status off* option.

Example 1

Suppresses the output of both parameters and the return status TDS token:

```
1> set proc_output_params OFF
2> go

1> set proc_return_status OFF
2> go

1> declare @x int
2> select @x = 2
3> exec sp_pout @x output
4> print "Value of @x returned from sproc is: %1!", @x
5> go

(1 row affected)
Value of @x returned from sproc is: 3
(1 row affected)
```


In addition, you can also suppress the lines reporting the number of rows affected to generate output with no extra messages using the `set nocount on` option before running this batch.

Example 2

Suppresses the output of parameter information:

```
1> create procedure sp_pout (@x int output) as select @x = @x + 1
2> go

1> set proc_output_params off
2> go

1> declare @x int
2> select @x = 1
3> exec sp_pout @x output
4> print "Value of @x returned from sproc is: %1!", @x
5> go
(1 row affected)
(return status = 0)
Value of @x returned from sproc is: 2
```

If you do not perform `set proc_output_params off`, the output after `(return status = 0)` includes the following:

Return parameters:

```
-----
                2
```

9.5 System Administration Guide

[CR #322843] In the *System Administration Guide*, Chapter 12, “Auditing,” Table 12-6, Values in event and extrainfo columns, is incomplete.

Table 12-6 should include these lines for the create index entry under:

“Other Information: Index name”

[CR #331120] In the *System Administration Guide*, Chapter 26, “Checking Database Consistency,” some size information is incorrect in the “Creating the dbccdb database” procedure. The procedure should read:

```
* In step 1, output:
Recommended size for dbccdb is 23MB (data = 21MB, log
= 2MB)
* In step 3, user input:
```

```
...
size = "4096"
...
size="1024"
* In step 7 user input:
use dbccdb|
go|
sp_dbcc_createws dbccdb, scanseg, scan_pubs2, scan,
"64K"|
sp_dbccvcreatews dbccdb, textseg, text_pubs2, text,
"64K"|
```

9.6 Performance and Tuning Guide

9.6.1 OAM - scans on DOL tables

OAM - pages are used when performing a table scan on a DOL-table. These OAM-pages refer to allocation-units, which must then be scanned to find the extents on the pages for the table being scanned. When the scan is started, a list of pages is generated the information in the allocation-page. Any newly allocated pages produced by another process in this allocation-unit are not seen by the scanning process as they were not allocated when the scan began. This is due to the way DOL-tables are implemented. If you want to include the newly allocated pages,

- use isolation level 3 type of scans (select with holdlock) or,
- add an index to the DOL-table which can be used to prevent a table-scan or,
- convert the DOL-table to an APL-table.

9.6.2 Multiple temp databases

The information on multiple tempdbs was inadvertently included in the *Performance and Tuning Guide*. This information should be in the *System Administration Guide* and will be moved there on the next update.

Login bindings to multiple tempdbs

In prior versions only a user with the “sa” role could be bound to a tempdb.

As of this version of Adaptive Server, any user can also be bound to an application at the time of login. For example:

If there is a binding for login “user1” and for application isql, when “user1” logs in with application isql the “user1” binding determines the tempdb assigned to the session, not the isql binding.

10. Technical support

Each Sybase installation that has purchased a support contract has one or more designated people who are authorized to contact Sybase Technical Support. If you have any questions about this installation or if you need assistance during the installation process, ask the designated person to contact Sybase Technical Support or the Sybase subsidiary in your area.

11. Other sources of information

Use the Sybase Getting Started CD, the Sybase Technical Library CD and the Technical Library Product Manuals Web site to learn more about your product:

- The Getting Started CD contains release bulletins and installation guides in PDF format, and may also contain other documents or updated information not included on the Technical Library CD. It is included with your software. To read or print documents on the Getting Started CD you need Adobe Acrobat Reader (downloadable at no charge from the Adobe Web site, using a link provided on the CD).
- The Technical Library CD contains product manuals and is included with your software. The DynaText reader (included on the Technical Library CD) allows you to access technical information about your product in an easy-to-use format.

Refer to the *Technical Library Installation Guide* in your documentation package for instructions on installing and starting the Technical Library.

- The Technical Library Product Manuals Web site is an HTML version of the Technical Library CD that you can access using a standard Web browser. In addition to product manuals, you will find links to EBFs/Updates, Technical Documents, Case Management, Solved Cases, newsgroups, and the Sybase Developer Network.

To access the Technical Library Product Manuals Web site, go to Product Manuals at <http://www.sybase.com/support/manuals/>.

11.1 Sybase certifications on the Web

Technical documentation at the Sybase Web site is updated frequently.

v **Finding the latest information on product certifications**

- 1 Point your Web browser to Technical Documents at <http://www.sybase.com/support/techdocs/>.
- 2 Select Products from the navigation bar on the left.
- 3 Select a product name from the product list and click Go.
- 4 Select the Certification Report filter, specify a time frame, and click Go.
- 5 Click a Certification Report title to display the report.

v **Creating a personalized view of the Sybase Web site (including support pages)**

Set up a MySybase profile. MySybase is a free service that allows you to create a personalized view of Sybase Web pages.

- 1 Point your Web browser to Technical Documents at <http://www.sybase.com/support/techdocs/>.
- 2 Click MySybase and create a MySybase profile.

11.2 Sybase EBFs and software maintenance

v **Finding the latest information on EBFs and software maintenance**

- 1 Point your Web browser to the Sybase Support Page at <http://www.sybase.com/support>.
- 2 Select EBFs/Maintenance. Enter user name and password information, if prompted (for existing Web accounts) or create a new account (a free service).
- 3 Select a product.
- 4 Specify a time frame and click Go.
- 5 Click the Info icon to display the EBF/Maintenance report, or click the product description to download the software.